

A generalized approach for Boolean matrix factorization

Rodrigo Cabral Farias^{a,*}, Sebastian Miron^b

^a *Université Côte d'Azur, CNRS, I3S Laboratory, 06900 Sophia-Antipolis, France*

^b *Université de Lorraine, CNRS, CRAN, 54000 Nancy, France*

Abstract

In this paper, we propose a generalized framework for fitting Boolean matrix factorization models to binary data. In this generalized setting, the binary rank-1 components of the underlying model can be combined by any Boolean function, thus extending the standard Boolean matrix factorization model, where the combination is restricted to the logical ‘OR’ function. We introduce two algorithms relying on a relaxation of the binary constraints on the factors of the model and on a polynomial representation of the Boolean function that combines the rank-1 components. One of the algorithms is based on the gradient descent optimization method, while the other is based on block coordinate descent. A detailed presentation of the algorithms is given, along with numerical experiments both on simulated and real datasets. A comparison with other algorithms from the literature is presented in the standard Boolean matrix setting allowing to assess the advantages and shortcomings of the proposed methods in terms of factor retrieval and data denoising performance, convergence behavior and time complexity.

Keywords: Binary data, Binary matrix factorizations, Boolean matrix factorizations, Data mining.

1. Introduction

Binary data matrices are one of the most natural ways of numerically encoding data in many applications. They can encode *yes/no* answers to surveys, voting records, tables indicating the presence or absence of traits of a group of individuals or indicating the proximity between the individuals, implicit feedback or thresholded explicit feedback (grades) in audio/video streaming platforms, input/output relationships in digital circuits, among many others. For this reason, a large number of data mining techniques specially tailored for binary data matrices have been developed

*Corresponding author

Email address: Rodrigo.CABRAL-FARIAS@univ-cotedazur.fr (Rodrigo Cabral Farias)

URL: <https://www.i3s.unice.fr/rcabral/en/home> (Rodrigo Cabral Farias)

8 recently. Among these techniques, a diverse number of models and algorithms relying on matrix
9 factorizations have appeared in the last two decades. They have been successfully used in a great
10 number of applications such as text mining [1], recommender systems [2–4], genetics [5, 6], protein
11 complex prediction [7], role mining [8] and telecommunications [9, 10].

12 Different factorization models for binary matrices have been proposed in the literature. One
13 group of models [2, 11–14] relies on a modification of logistic regression where a logistic function is
14 applied to a constrained matrix factorization model such as principal component analysis (PCA).
15 Most models of this group are named logistic PCA or binary PCA. Although the matrix factors
16 are not constrained to be binary, the use of the logistic function allows the model to constrain the
17 elements of its output matrix to lie in the interval $[0, 1]$. A generalized version of logistic PCA,
18 based on the family of mean-parameterized Bernoulli models, was recently studied in [15].

19 Another model, called binary matrix factorization (bMF) [16], directly factorizes the data matrix
20 into two binary matrices. This model is equivalent to a decomposition of the data matrix into a
21 sum of rank-1 binary matrices. A major issue for fitting optimally a bMF model to data is the
22 discrete nature of the model parameters, which makes the underlying optimization problem difficult
23 to be solved. It has been shown that fitting a single component bMF model is already a NP-hard
24 problem [17]. Due to its hardness, algorithms for bMF do not aim to solve exactly the original fitting
25 problem. A class of methods, such as the association rules algorithm (ASSO) [18] or formal concept
26 analysis (FC) [19], rely on low complexity iterative rules that extract in a greedy-like manner binary
27 rank-1 components that are expected to approximate the optimal ones. A revisited version of the
28 FC-based algorithms of [19], significantly faster, was recently introduced in [20]; another variant,
29 that uses the minimum description length principle (MDL) for factor selection was proposed in [21].
30 A different approach, called the penalty function algorithm (PF) [16], solves a relaxed version of the
31 underlying fitting problem. In PF, the bMF factor elements are relaxed to the nonnegative orthant,
32 while a penalization term forcing the factors to be close to binary is added to the original data fitting
33 objective function. Such a relaxation allows to use multiplicative gradient algorithms to retrieve the
34 factors in a similar manner as for nonnegative matrix factorization (NMF) [22].

35 The limitations of bMF appear whenever its expected rank-1 components have overlapping sup-
36 ports. In this case, the sum of the rank-1 components does not lead to a binary matrix. One
37 way to counter this issue is to assume that the presence of a ‘1’ in the data matrix is due to the
38 contributions of several ‘1’ in the rank-1 terms. Note that this corresponds to simply replacing the
39 arithmetic sums in the decomposition model by logical ‘OR’ operations. This leads to a particular

40 matrix factorization model, called Boolean matrix factorization (BMF).

41 A modification of the PF algorithm explicitly tailored for BMF has been proposed in [23]. The
42 authors propose to apply a threshold function to the bMF model, such that the output matrix
43 elements are either 0 or 1. Since the threshold function is not differentiable, to be able to use a
44 multiplicative gradient algorithm as in the PF algorithm, a smooth approximation of the threshold
45 function is used. The resulting BMF method is called post-nonlinear PF algorithm (PNL-PF).
46 A heuristic model selection algorithm for estimating the number of binary sources in the BMF
47 setting was also proposed in [24], based on stability criteria. This method constructs an ensemble of
48 random matrices that are slight perturbations of the initial matrix to test the stability of the Boolean
49 decomposition. Other algorithms have also been developed under a stochastic setting, where the
50 elements of the factors are supposed to be random [25, 26].

51 In this paper, we propose an approximate factorization approach for binary valued matrices that
52 generalizes BMF to arbitrary Boolean “*sum*” functions. Instead of considering combinations of the
53 rank-1 components with logical ‘OR’, we assume that an arbitrary Boolean function with known
54 truth table is used. Our approach is based on the relaxation of the binary constraints, as in PF and
55 PNL-PF, but instead of representing the behavior of the logical combiner with a threshold function,
56 we represent it as a multivariate polynomial of the elements of each component. Since a multivariate
57 polynomial is a differentiable function, such a representation allows developing a gradient algorithm
58 for fitting the generalized BMF model, without the need to resort to smooth approximations, as in
59 PNL-PF. We also propose a generalized BMF approximation algorithm based on block coordinate
60 descent. The algorithm alternatively updates the columns of the factors to be retrieved in a similar
61 manner as in hierarchical alternating least squares (HALS) for NMF [27]. Since the multivariate
62 polynomial representing the logical combiner is multilinear in the elements of the components, the
63 updates required in the block coordinate descent algorithm can be obtained in closed-form.

64 We present implementation details of our approach in the specific case of of BMF approximation,
65 and under this setting we compare the performance of the two resulting algorithms with state-of-the
66 art BMF methods. The performance of the methods are evaluated in terms of denoising and factor
67 retrieval capabilities, but also in terms of convergence behavior, time complexity and sensitivity to
68 initializations. To illustrate our approach in a more practical setting, we apply one of the proposed
69 methods to retrieve the BMF of 4 real datasets. Finally, we also show simulation results concerning
70 the application of the general version of the proposed algorithms to retrieve factorizations where
71 the component combining functions are the logical exclusive ‘OR’ (‘XOR’) and the 3-term majority

72 function.

73 1.1. Outline

74 In Section 2, we present the binary and Boolean factorization models along with the polynomial
75 representation of the general Boolean factorization. We introduce two algorithms for the generalized
76 Boolean factorization in Section 3 and illustrate their implementation in the specific case of BMF.
77 Section 4 shows the results of the conducted numerical experiments to compare the performance of
78 the proposed algorithms with state-of-the art methods. Results on 4 real datasets are also given.
79 We also provide numerical simulation results in a more general setting, where the data do not follow
80 the standard BMF model. Finally, we conclude the paper in Section 5.

81 1.2. Notations

82 Scalars are represented by lower case letters x , while vectors are represented by bold-face lower
83 case letters \mathbf{x} . Bold-face upper case letters \mathbf{X} are used to represent matrices. A single subscript x_i
84 is used to represent the i -th element of a vector or the i -th column of a matrix \mathbf{x}_i (i -th column of
85 \mathbf{X}). A double subscript x_{ij} denotes the (i, j) -th element of a matrix. Superscripts (or subscripts)
86 of the form $\mathbf{x}^{1:n}$ denote the tuple $(\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n)$.

87 Matrix transpose is denoted \mathbf{X}^T , while the Frobenius norm of a matrix is symbolized by $\|\mathbf{X}\|_F$.
88 To denote a matrix of size $I \times J$ with all elements equal to 1, we use $\mathbf{1}_{I \times J}$. The symbol \square denotes
89 the Hadamard (entry-wise) matrix product and $\text{Diag}(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n)$ denotes a block diagonal
90 matrix with matrices $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$ in its diagonal blocks. The column-major vectorization of
91 a matrix is denoted $\text{vec}(\mathbf{X})$.

92 Logical ‘OR’ is symbolized by \vee and the same symbol is used for its entry-wise matrix version.
93 Logical ‘XOR’ is denoted by \oplus .

94 2. General binary and Boolean factorizations

95 We are interested in exactly or approximately decomposing a $I \times J$ data matrix \mathbf{Y} with binary
96 elements $y_{ij} \in \{0, 1\}$, for $(i, j) \in \{1, 2, \dots, I\} \times \{1, 2, \dots, J\}$, into $R \geq 2$ binary rank-1 matrices
97 $\mathbf{X}^{1:R} = \{\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^R\}$. Each binary \mathbf{X}^r with $r \in \{1, 2, \dots, R\}$ is written as

$$\mathbf{X}^r = \mathbf{a}_r \mathbf{b}_r^T, \quad (1)$$

98 where \mathbf{a}_r and \mathbf{b}_r are vectors of sizes I and J respectively and with their elements constrained
99 to be binary $[\mathbf{a}_r]_i \in \{0, 1\}$, for $(i, r) \in \{1, 2, \dots, I\} \times \{1, 2, \dots, R\}$, $[\mathbf{b}_r]_j \in \{0, 1\}$, for $(j, r) \in$

100 $\{1, 2, \dots, J\} \times \{1, 2, \dots, R\}$. The vectors \mathbf{a}_r and \mathbf{b}_r can be stored in matrices $\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_R]$
 101 and $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_R]$. As presented in [23] and [28], different decompositions can be considered
 102 depending on how one precisely defines the way that \mathbf{X}_r are combined to approximate \mathbf{Y} . If we
 103 specify a function $f : \{0, 1\}^R \rightarrow \mathcal{Y} \subset \mathbb{R}$ defined on R binary inputs and resulting in a value on a
 104 finite subset \mathcal{R} of the integers, general binary factorization corresponds to approximate the elements
 105 of \mathbf{Y} as follows:

$$y_{ij} \approx f(x_{ij}^1, x_{ij}^2, \dots, x_{ij}^R) = f(x_{ij}^{1:R}) = f(a_{i,1:R} b_{j,1:R}), \quad (2)$$

106 where $a_{i,1:R} b_{j,1:R} = \{a_{i1} b_{j1}, a_{i2} b_{j2}, \dots, a_{iR} b_{jR}\}$. Denoting the matrix resulting of the element-wise
 107 application of $f(\cdot)$ to the rank-one matrices $\mathbf{X}^{1:R}$ simply by $f(\mathbf{X}^{1:R})$, the approximation problem
 108 (2) can be cast as the following minimization problem:

$$\begin{aligned} & \text{minimize} && \mathcal{F}(\mathbf{A}, \mathbf{B}) = \frac{1}{2} \|\mathbf{Y} - f(\mathbf{X}^{1:R})\|_F^2 \\ & \text{(where} && \mathbf{X}^{1:R} = \{\mathbf{X}^1 = \mathbf{a}_1 \mathbf{b}_1^\top, \dots, \mathbf{X}^R = \mathbf{a}_R \mathbf{b}_R^\top\}, \\ & \text{with respect to} && \mathbf{A} \in \{0, 1\}^{I \times R}, \mathbf{B} \in \{0, 1\}^{J \times R}. \end{aligned} \quad (3)$$

109 A solution for this problem is guaranteed to exist since the cardinal of the feasible set is finite.
 110 If a solution $\mathbf{A}^*, \mathbf{B}^*$ of (3) achieves $\mathcal{F}(\mathbf{A}, \mathbf{B}) = 0$, we say that it is an exact factorization of \mathbf{Y} . In
 111 the data analysis literature, 3 common types of factorization problems which are special cases of the
 112 general form above are the following:

113 *Binary matrix factorization (bMF)*. When $f(x_{ij}^{1:R}) = \sum_{r=1}^R x_{ij}^{1:R}$ is the usual sum on \mathbb{R} . See the left
 114 column of Tab. 1 for an example when $R = 2$. We say that \mathbf{A} and \mathbf{B} are approximate factors of \mathbf{Y} ,
 115 since in this case $\mathbf{Y} \approx \mathbf{A}\mathbf{B}^\top$.

116 *Boolean matrix factorization (BMF)*. When $f(x_{ij}^{1:R}) = \bigvee_{r=1}^R x_{ij}^{1:R}$ is the R -term logical ‘OR’. See
 117 the middle column of Tab. 1 for an example. Similarly to the previous case, we can say that \mathbf{A} and
 118 \mathbf{B} are approximate Boolean factors of \mathbf{Y} , since $\mathbf{Y} \approx \mathbf{A} \wedge \mathbf{B}^\top$ where $(\cdot \wedge \cdot)$ is the matrix product
 119 defined in the Boolean semi-ring (sums are replaced by logical ‘OR’).

120 \mathbb{F}_2 *matrix factorization (F2MF)*. When $f(x_{ij}^{1:R}) = \bigoplus_{r=1}^R x_{ij}^{1:R}$ is the R -term modulo-2 sum, that
 121 is, a cascade of R logical ‘XOR’ operations applied sequentially to $x_{ij}^{1:R}$. In this case, we can write
 122 $\mathbf{Y} \approx \mathbf{A} \odot \mathbf{B}^\top$, where $(\cdot \odot \cdot)$ is the matrix product in the \mathbb{F}_2 field (Galois field of two elements). Here
 123 the sums are replaced by logical ‘XOR’. Therefore, we call this model \mathbb{F}_2 matrix factorization.

124 Observe that if one wants to factorize a binary data matrix \mathbf{Y} without errors using bMF, its
 125 factors should contain columns with disjoint supports. This is due to the presence of possible values
 126 larger than 1 in the outputs of the sum operation for bMF (see Tab. 1).

Inputs $x^{(1)}, x^{(2)}$	Output		
	bMF (+)	BMF (\vee)	F2MF (\oplus)
0, 0	0	0	0
0, 1	1	1	1
1, 0	1	1	1
1, 1	2	1	0

Table 1: Results for different $f(\cdot)$ with $R = 2$ inputs used in different factorizations.

127 For a given \mathbf{Y} , the characteristics of its factorization such as rank or uniqueness may change
128 depending on the chosen $f(\cdot)$. Before focusing on algorithms for general Boolean factorizations,
129 which are the main contribution of this paper, we briefly illustrate with some toy examples, some
130 important differences between factorizations with different $f(\cdot)$.

131 2.1. Ranks and uniqueness of binary factorizations

132 As in standard matrix factorizations, the minimal number of columns R for which exact factor-
133 izations of \mathbf{Y} exist is called the rank of \mathbf{Y} and we denote it $\text{rank}_f(\mathbf{Y})$:

$$\text{rank}_f(\mathbf{Y}) = \min \left\{ R \mid \mathbf{Y} = f(\mathbf{X}^{1:R}), \mathbf{A} \in \{0, 1\}^{I \times R}, \mathbf{B} \in \{0, 1\}^{J \times R} \right\}. \quad (4)$$

134 Following the denominations in [23, 28], if $f(\cdot)$ is the standard sum, we call this rank the *binary rank*
135 and we denote it $\text{rank}_{\{0,1\}}(\mathbf{Y})$. If $f(\cdot)$ is the logical ‘OR’ then this rank is the *Boolean rank* and it
136 is denoted $\text{rank}_{\mathbb{B}}(\mathbf{Y})$. We call it \mathbb{F}_2 *rank*, when the combining function is the modulo-2 sum and we
137 denote it $\text{rank}_{\mathbb{F}_2}(\mathbf{Y})$. In what follows, we give some toy examples from the literature to illustrate
138 the fact that these ranks can be different for a given matrix. Consider the 3×3 binary matrix [18]

$$\mathbf{Y}_1 = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix}. \quad (5)$$

139 Minimum rank decompositions of \mathbf{Y}_1 for bMF and BMF are

$$\mathbf{Y}_1 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}^T + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}^T + \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}^T = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}^T \vee \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}^T$$

140 and the same factors of bMF can be used for F2MF. Thus, we have $\text{rank}_{\mathbb{F}_2}(\mathbf{Y}_1) = 3$, which is
141 larger than $\text{rank}_{\mathbb{B}}(\mathbf{Y}_1) = 2$. Note also that the rank of \mathbf{Y}_1 on the reals is $\text{rank}(\mathbf{Y}_1) = 3$, since the 3
142 columns of \mathbf{Y}_1 are linearly independent.

143 One can easily find cases where the relations between these ranks are different from the previous
144 example. Consider a matrix \mathbf{Y}_2 which is equal to \mathbf{Y}_1 except for the central element $[Y_1]_{2,2}$ which is

145 flipped to zero. Then the BMF factors for \mathbf{Y}_1 give an exact F2MF for \mathbf{Y}_2 with a minimum number
 146 of columns. In this case, $\text{rank}_{\mathbb{B}}(\mathbf{Y}_2) = \text{rank}(\mathbf{Y}_2) = 3$, but $\text{rank}_{\mathbb{F}_2}(\mathbf{Y}_2) = 2$.

147 By increasing the size of the data matrix, one can also find cases where the $\text{rank}(\mathbf{Y}) < \text{rank}_{\{0,1\}}(\mathbf{Y})$.
 148 For example, for [29]

$$\mathbf{Y}_3 = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix},$$

149 we have $\text{rank}(\mathbf{Y}_3) = 3$, while $\text{rank}_{\{0,1\}}(\mathbf{Y}_3) = \text{rank}_{\mathbb{B}}(\mathbf{Y}_3) = 4$.

150 When the rank-one components \mathbf{X}^r have disjoint supports, all of the previously mentioned ranks
 151 coincide. Thus, for [23]

$$\mathbf{Y}_4 = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \quad (6)$$

152 $\text{rank}_{\{0,1\}}(\mathbf{Y}_4) = \text{rank}_{\mathbb{B}}(\mathbf{Y}_4) = \text{rank}_{\mathbb{F}_2}(\mathbf{Y}_4) = \text{rank}(\mathbf{Y}_4) = 2$, and one exact decomposition corre-
 153 sponding to this rank has factors

$$\mathbf{A} = \mathbf{B} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{bmatrix}, \quad (7)$$

154 for all factorizations presented above, but also for factorizations on \mathbb{R} .

155 In data mining applications, for interpretability reasons, it is expected that factors \mathbf{A} and \mathbf{B}
 156 could be retrieved uniquely from the data up to joint column permutations. It is well-known that,
 157 in general, matrix factorization over the real numbers is not unique. For F2MF, the factorization
 158 is unique only when $R = 1$, since for $R \geq 2$ one can find binary \mathbf{T}, \mathbf{T}' , different from permutation
 159 matrices, such that $\mathbf{A} \odot \mathbf{B}^{\top} = (\mathbf{A} \odot \mathbf{T}) \odot (\mathbf{B} \odot \mathbf{T}')^{\top}$. Regarding bMF and BMF, the binary constraints
 160 on \mathbf{A} and \mathbf{B} allow to retrieve unique factors under some particular conditions (see [23, 28, 30–32]
 161 for details on uniqueness conditions).

162 2.2. Polynomial representation of a general Boolean function

163 We focus in this paper in solving problem (3) where $f(\cdot)$ is a general Boolean function $f :$
 164 $\{0, 1\}^R \rightarrow \{0, 1\}$.

165 We consider a two-step approach: in the first step, we apply an optimization algorithm to solve
 166 a relaxed version of (3) where the binary constraints are dropped. The elements of \mathbf{A} and \mathbf{B} are

167 either allowed to lie on \mathbb{R} or on the interval $[0, 1]$. In the second step, the resulting approximations
 168 of \mathbf{A} and \mathbf{B} , denoted $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$, are projected onto the set of binary values. This projection $\mathcal{P}_{\mathbb{B}}(\cdot)$
 169 corresponds simply to a thresholding operation. For example, for the elements of $\hat{\mathbf{A}}$:

$$\left[\mathcal{P}_{\mathbb{B}}(\hat{\mathbf{A}})\right]_{ir} = \begin{cases} 0 & , \text{ for } \hat{a}_{ir} < 0.5, \\ 1 & , \text{ for } \hat{a}_{ir} \geq 0.5. \end{cases} \quad (8)$$

170 To apply this approach, the Boolean function $f(\cdot)$ must be represented by another function \bar{f} defined
 171 for real inputs. \bar{f} should be defined in such a way that both functions are equal for binary inputs.
 172 In this work, we define $\bar{f}(\cdot)$ relying on the fact that any Boolean function of R variables $\mathbf{x} =$
 173 $[x_1, x_2, \dots, x_R]^T$ can be written as a multivariate polynomial. For any $\mathbf{x} \in \{0, 1\}^R$, the following
 174 polynomial achieves the same values as $f(\cdot)$ [33]:

$$\bar{f}(\mathbf{x}) = \sum_{\mathbf{w} \in \mathcal{W}^1} \left\{ \prod_{i|w_i=1} x_i \prod_{j|w_j=0} (1-x_j) \right\}, \quad (9)$$

175 where $\mathcal{W}^1 = \{\mathbf{w} \in \{0, 1\}^R \mid f(\mathbf{w}) = 1\}$.

176 Examples of polynomial representation of simple Boolean functions are the following:

- 177 • Logical ‘OR’ with $R = 2$, $(x_1 \vee x_2)$:

$$\bar{f}_{\text{OR}}(x_1, x_2) = (1-x_1)x_2 + x_1(1-x_2) + x_1x_2. \quad (10)$$

- 178 • Logical ‘XOR’ with $R = 2$, $(x_1 \oplus x_2)$:

$$\bar{f}_{\text{XOR}}(x_1, x_2) = (1-x_1)x_2 + x_1(1-x_2). \quad (11)$$

- 3-term majority:

$$\begin{aligned} \bar{f}_{\text{MAJ}}(x_1, x_2, x_3) = \mathbf{1}_{(\sum_i x_i) \geq 2}(x_1, x_2, x_3) &= (1-x_1)x_2x_3 + x_1(1-x_2)x_3 \\ &+ x_1x_2(1-x_3) + x_1x_2x_3. \end{aligned} \quad (12)$$

179 Observe that with this representation, the number of terms in the polynomial depends on the car-
 180 dinal of \mathcal{W}^1 (number of input combinations such that $f(\mathbf{x}) = 1$). If the set $\mathcal{W}^0 = \{\mathbf{w} \in \{0, 1\}^R \mid f(\mathbf{w}) = 0\}$
 181 has a smaller cardinal than \mathcal{W}^1 , then it may be more convenient to use another equivalent form of
 182 $\bar{f}(\cdot)$:

$$\bar{f}(\mathbf{x}) = 1 - \sum_{\mathbf{w} \in \mathcal{W}^0} \left\{ \prod_{i|w_i=1} x_i \prod_{j|w_j=0} (1-x_j) \right\}. \quad (13)$$

183 Note that in the case of $x_1 \vee x_2$, the representation above leads to $\bar{f}_{\text{OR}}(x_1, x_2) = 1 - (1 - x_1)(1 - x_2)$
 184 and the corresponding R -term version of ‘OR’ has a simple expression:

$$\bar{f}_{\text{OR}}(x_1, x_2, \dots, x_R) = 1 - \prod_{r=1}^R (1 - x_r). \quad (14)$$

185 In the rest of the paper, we use representation (13), since it allows an easier presentation of the
 186 algorithms that we propose in the specific case of BMF.

187 3. Algorithms

188 Two properties of $\bar{f}(\cdot)$ are interesting from an optimization point of view: this function is differ-
 189 entiable and it is multilinear in its inputs. Since $\bar{f}(\cdot)$ is differentiable, gradient descent can be applied
 190 to attempt solving the corresponding relaxed versions of (3). Multilinearity of $\bar{f}(\cdot)$ with respect to
 191 its inputs implies that this function is also multilinear in the columns of \mathbf{A} and \mathbf{B} . Therefore, if we
 192 use a block-coordinate descent approach to attempt minimizing relaxed (3), and we set the blocks
 193 of variables to be the columns of \mathbf{A} and \mathbf{B} , the block updates will be given by the solutions of
 194 simple linear least squares problems. As a consequence, each of these properties leads to a different
 195 algorithm for solving relaxed (3). These algorithms are detailed next.

196 3.1. Gradient descent (GD) algorithm

197 In the first algorithm, we consider a relaxed version of (3) where elements of the factors are
 198 allowed to lie in \mathbb{R} . To force the solution to be close to binary, we introduce a penalty term $\mathcal{G}(\mathbf{A}, \mathbf{B})$
 199 in the objective function as in [16, 23]. The expression of this penalty term is

$$\mathcal{G}(\mathbf{A}, \mathbf{B}) = \frac{1}{2} \left\{ \sum_{i=1}^I \sum_{j=1}^J \sum_{r=1}^R [a_{ir}^2 (1 - a_{ir})^2 + b_{jr}^2 (1 - b_{jr})^2] \right\}. \quad (15)$$

200 Note that this penalty is minimal whenever all elements of the factors are ‘0’ or ‘1’. The new
 201 optimization problem we have to solve is the following:

$$\begin{aligned} & \text{minimize} && \mathcal{H}(\mathbf{A}, \mathbf{B}; \lambda) = \mathcal{F}(\mathbf{A}, \mathbf{B}) + \lambda \mathcal{G}(\mathbf{A}, \mathbf{B}) \\ & \text{with respect to} && \mathbf{A} \in \mathbb{R}^{I \times R}, \mathbf{B} \in \mathbb{R}^{J \times R}, \end{aligned} \quad (16)$$

202 where $\lambda > 0$ is a given value. Since this penalty term is differentiable, $\mathcal{H}(\mathbf{A}, \mathbf{B}; \lambda)$ is differentiable.
 203 Therefore, we can apply the standard gradient descent algorithm to find its critical points.

204 In standard gradient descent, the entries of the parameters vector $\boldsymbol{\theta} = [\text{vec}(\mathbf{A})^\top \text{vec}(\mathbf{B})^\top]^\top$,
 205 where $\text{vec}(\mathbf{A}) = [\mathbf{a}_1^\top, \mathbf{a}_2^\top, \dots, \mathbf{a}_R^\top]^\top$ and $\text{vec}(\mathbf{B}) = [\mathbf{b}_1^\top, \mathbf{b}_2^\top, \dots, \mathbf{b}_R^\top]^\top$, are estimated jointly. The

206 estimate $\hat{\boldsymbol{\theta}}_k$ of the parameter vector at iterate k is given by

$$\hat{\boldsymbol{\theta}}_k = \hat{\boldsymbol{\theta}}_{k-1} - \gamma_k \nabla_{\boldsymbol{\theta}} \mathcal{H}(\boldsymbol{\theta})|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}_{k-1}}, \quad (17)$$

207 where γ_k is the step-size of the algorithm and $\nabla_{\boldsymbol{\theta}} \mathcal{H}(\boldsymbol{\theta})|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}_{k-1}}$ is the gradient vector of $\mathcal{H}(\cdot)$ with
 208 respect to all parameters $\boldsymbol{\theta}$ evaluated at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}_{k-1}$.

209 *Gradient expressions.* The full gradient vector can be written as a function of the partial gradients
 210 with respect to \mathbf{A} and \mathbf{B} as follows

$$\nabla_{\boldsymbol{\theta}}^{\top} \mathcal{H}(\boldsymbol{\theta}) = \left[\nabla_{\text{vec}(\mathbf{A})}^{\top} \mathcal{H}(\mathbf{A}, \mathbf{B}) \quad \nabla_{\text{vec}(\mathbf{B})}^{\top} \mathcal{H}(\mathbf{A}, \mathbf{B}) \right] \quad (18)$$

211 and the partial gradients are

$$\nabla_{\text{vec}(\mathbf{A})}^{\top} \mathcal{H}(\mathbf{A}, \mathbf{B}) = \left[\frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial a_{11}} \quad \dots \quad \frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial a_{I1}} \quad \dots \quad \frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial a_{1R}} \quad \dots \quad \frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial a_{IR}} \right], \quad (19)$$

$$\nabla_{\text{vec}(\mathbf{B})}^{\top} \mathcal{H}(\mathbf{A}, \mathbf{B}) = \left[\frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial b_{11}} \quad \dots \quad \frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial b_{J1}} \quad \dots \quad \frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial b_{1R}} \quad \dots \quad \frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial b_{JR}} \right]. \quad (20)$$

213 The elements of $\nabla_{\text{vec}(\mathbf{A})} \mathcal{H}(\mathbf{A}, \mathbf{B})$ are given by

$$\frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial a_{i'r'}} = - \left\{ \sum_{j=1}^J \left[y_{i'j} - \bar{f}(x_{i'j}^{1:R}) \right] \frac{\partial \bar{f}(x_{i'j}^{1:R})}{\partial x_{i'j}^{r'}} b_{j'r'} \right\} + \lambda a_{i'r'} (1 - a_{i'r'}) (1 - 2a_{i'r'}) \quad (21)$$

for $i' \in \{1, \dots, I\}$ and $r' \in \{1, \dots, R\}$, where the expression of the partial derivatives of $\bar{f}(\cdot)$ are

$$\begin{aligned} \frac{\partial \bar{f}(x_{i'j}^{1:R})}{\partial x_{i'j}^{r'}} &= \sum_{\substack{\mathbf{w} \in \mathcal{W}^0, \\ w_{r'} = 0}} \left[\prod_{s|w_s=1} x_{i'j}^{s'} \right] \left[\prod_{\substack{s'|w'_s=0, \\ s' \neq r'}} (1 - x_{i'j}^{s'}) \right] \\ &\quad - \sum_{\substack{\mathbf{w} \in \mathcal{W}^0, \\ w_{r'} = 1}} \left[\prod_{\substack{s|w_s=1, \\ s \neq r'}} x_{i'j}^s \right] \left[\prod_{s'|w'_s=0} (1 - x_{i'j}^{s'}) \right]. \end{aligned} \quad (22)$$

214 The (j', r') element of $\nabla_{\text{vec}(\mathbf{B})} \mathcal{H}(\mathbf{A}, \mathbf{B})$ for $j' \in \{1, \dots, J\}$, $r' \in \{1, \dots, R\}$ is

$$\frac{\partial \mathcal{H}(\mathbf{A}, \mathbf{B})}{\partial b_{j'r'}} = - \left\{ \sum_{i=1}^I \left[y_{ij'} - \bar{f}(x_{ij'}^{1:R}) \right] \frac{\partial \bar{f}(x_{ij'}^{1:R})}{\partial x_{ij'}^{r'}} a_{i,r'} \right\} + \lambda b_{j'r'} (1 - b_{j'r'}) (1 - 2b_{j'r'}). \quad (23)$$

The partial gradients can be written in vector form as a function of \mathbf{A} and \mathbf{B} as follows

$$\begin{aligned} \nabla_{\text{vec}(\mathbf{A})} \mathcal{H}(\mathbf{A}, \mathbf{B}) &= -\text{Diag}(\mathbf{E} \square \mathbf{P}_1, \dots, \mathbf{E} \square \mathbf{P}_R) \text{vec}(\mathbf{B}) \\ &\quad + \lambda \text{vec}(\mathbf{A}) \square (\mathbf{1}_{IR \times 1} - \text{vec}(\mathbf{A})) \square (\mathbf{1}_{IR \times 1} - 2\text{vec}(\mathbf{A})), \\ \nabla_{\text{vec}(\mathbf{B})} \mathcal{H}(\mathbf{A}, \mathbf{B}) &= -\text{Diag}(\mathbf{E}^{\top} \square \mathbf{P}_1^{\top}, \dots, \mathbf{E}^{\top} \square \mathbf{P}_R^{\top}) \text{vec}(\mathbf{A}) \\ &\quad + \lambda \text{vec}(\mathbf{B}) \square (\mathbf{1}_{JR \times 1} - \text{vec}(\mathbf{B})) \square (\mathbf{1}_{JR \times 1} - 2\text{vec}(\mathbf{B})), \end{aligned} \quad (24)$$

215 where \mathbf{E} is the model error matrix

$$\mathbf{E} = \mathbf{Y} - \bar{f}(\mathbf{X}^{1:R}) = \mathbf{Y} - \mathbf{1}_{I \times J} + \sum_{\mathbf{w} \in \mathcal{W}^0} \left[\begin{array}{c} \square \\ \square \end{array} \begin{array}{c} \mathbf{X}^s \\ \mathbf{X}^s \end{array} \right] \square \left[\begin{array}{c} \square \\ \square \end{array} \begin{array}{c} (\mathbf{1}_{I \times J} - \mathbf{X}^{s'}) \\ (\mathbf{1}_{I \times J} - \mathbf{X}^{s'}) \end{array} \right] \quad (25)$$

and $\mathbf{P}_{r'}$ are $I \times J$ matrices given by

$$\begin{aligned} \mathbf{P}_{r'} &= \sum_{\substack{\mathbf{w} \in \mathcal{W}^0, \\ w_{r'} = 0}} \left[\begin{array}{c} \square \\ \square \end{array} \begin{array}{c} \mathbf{X}^s \\ \mathbf{X}^s \end{array} \right] \square \left[\begin{array}{c} \square \\ \square \end{array} \begin{array}{c} (\mathbf{1}_{I \times J} - \mathbf{X}^{s'}) \\ (\mathbf{1}_{I \times J} - \mathbf{X}^{s'}) \end{array} \right] \\ &\quad \substack{s|w_s=1, \\ s' \neq r'} \\ &- \sum_{\substack{\mathbf{w} \in \mathcal{W}^0, \\ w_{r'} = 1}} \left[\begin{array}{c} \square \\ \square \end{array} \begin{array}{c} \mathbf{X}^s \\ \mathbf{X}^s \end{array} \right] \square \left[\begin{array}{c} \square \\ \square \end{array} \begin{array}{c} (\mathbf{1}_{I \times J} - \mathbf{X}^{s'}) \\ (\mathbf{1}_{I \times J} - \mathbf{X}^{s'}) \end{array} \right]. \end{aligned} \quad (26)$$

216 *Step-size, penalty constant and initializations.* In the simplest version of the algorithm the step-size
 217 γ_k can be set to a small constant value. The penalty factor λ may be chosen as variable through
 218 iterations: λ is set to a value close to zero in the first iterations and its increased up to a high target
 219 value.

220 Since the cost function being minimized is highly nonconvex, gradient descent may converge
 221 to spurious critical points. For this reason, it is important to test different initializations of the
 222 algorithm and pick the solution which gives the best data fitting. The algorithm can be initialized
 223 each time with different random elements for the factor updates $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$. The elements of $\hat{\mathbf{A}}$
 224 and $\hat{\mathbf{B}}$ can be drawn from independent and identically distributed (iid) uniform samples: $\hat{a}_{ir} \sim$
 225 $\mathcal{U}[0, 1]$, $\hat{b}_{jr} \sim \mathcal{U}[0, 1]$, for $i \in \{1, \dots, I\}$, $j \in \{1, \dots, J\}$ and $r \in \{1, \dots, R\}$.

226 3.2. Projected Hierarchical Alternating Least Squares (PHALS) algorithm

227 In the second approach, named projected hierarchical alternating least squares (PHALS), the
 228 cost function is minimized with respect to each $\mathbf{a}_1, \dots, \mathbf{a}_R, \mathbf{b}_1, \dots, \mathbf{b}_R$ in an alternating manner,
 229 similar to the hierarchical alternating least squares (HALS) method [27]. The minimization with
 230 respect to each column is performed by relaxing the binary constraints to $\mathbb{R}^I, \mathbb{R}^J$. After updating
 231 all the estimates of a column of a factor, we project elements of the updated factor onto the interval
 232 $[0, 1]$ to prevent the updates to converge to negative or large positive values.

Suppose that we want to update the estimate $\hat{\mathbf{a}}_{r'}$ of $\mathbf{a}_{r'}$, all other columns of the factors are then
 considered to be equal to $\hat{\mathbf{a}}_r$ with $r \neq r'$ and $\hat{\mathbf{b}}_r$ for $r \in \{1, \dots, R\}$. The updated $\hat{\mathbf{a}}_{r'}$ is then given

by the minimization of

$$\begin{aligned}\mathcal{F}_{\text{PHALS}}(\mathbf{a}_{r'}) &= \frac{1}{2} \sum_{i=1}^I \sum_{j=1}^J \left\{ y_{ij} - 1 + a_{ir'} \hat{b}_{jr'} \hat{p}_{ij}^{r'} + (1 - a_{ir'} \hat{b}_{jr'}) \hat{q}_{ij}^{r'} \right\}^2 \\ &= \frac{1}{2} \sum_{i=1}^I \sum_{j=1}^J \left[y_{ij} - 1 + \hat{q}_{ij}^{r'} + a_{ir'} \hat{b}_{jr'} (\hat{p}_{ij}^{r'} - \hat{q}_{ij}^{r'}) \right]^2,\end{aligned}\quad (27)$$

233 where

$$\hat{p}_{ij}^{r'} = \sum_{\substack{\mathbf{w} \in \mathcal{W}^0, \\ w_{r'} = 1}} \left[\prod_{\substack{s | w_s = 1 \\ s \neq r'}} \hat{a}_{is} \hat{b}_{js} \right] \left[\prod_{\substack{s' | w_{s'} = 0 \\ s' \neq r'}} (1 - \hat{a}_{is} \hat{b}_{js}) \right] \quad (28)$$

234 and $\hat{q}_{ij}^{r'}$ is similarly defined except that the summation is done through $\mathbf{w} \in \mathcal{W}^0$ whose $w_{r'} = 0$.

235 The cost function $\mathcal{F}_{\text{PHALS}}(\mathbf{a}_{r'})$ can be rewritten as $\mathcal{F}_{\text{PHALS}}(\mathbf{a}_{r'}) = \sum_{i=1}^I \mathcal{F}_i(a_{ir'})$, where, for a
236 given i' ,

$$\mathcal{F}_{i'}(a_{i'r'}) = \sum_{j=1}^J \left(y_{i'j} - 1 + \hat{q}_{i'j}^{r'} + a_{i'r'} \hat{b}_{jr'} (\hat{p}_{i'j}^{r'} - \hat{q}_{i'j}^{r'}) \right)^2. \quad (29)$$

237 Observe that each term $\mathcal{F}_i(a_{ir'})$ of the cost function depends only on one of the $a_{ir'}$, thus the elements
238 of $\hat{\mathbf{A}}_{r'}$ can be obtained separately by minimizing $\mathcal{F}_{i'}(a_{i'r'})$. The function $\mathcal{F}_{i'}(a_{i'r'})$ is quadratic on
239 $a_{i'r'}$, therefore its unconstrained minimum can be easily obtained. For a given i' , it is

$$\hat{a}_{i'r'} = \frac{\sum_{j=1}^J \left[(y_{i'j} - 1 + \hat{q}_{i'j}^{r'}) (\hat{q}_{i'j}^{r'} - \hat{p}_{i'j}^{r'}) \hat{b}_{jr'} \right]}{\sum_{j=1}^J \left[\hat{b}_{jr'} (\hat{q}_{i'j}^{r'} - \hat{p}_{i'j}^{r'}) \right]^2}. \quad (30)$$

240 Once $\hat{\mathbf{a}}_{r'}$ has been completely updated, the arrays $\hat{p}_{i,j}^r$ and $\hat{q}_{i,j}^r$ have to be recalculated for the
241 update of the next $\hat{\mathbf{a}}_r$. When all columns of $\hat{\mathbf{A}}$ have been updated, the projection of its elements
242 onto $[0, 1]$ is given by

$$\left[\mathcal{P}_U(\hat{\mathbf{A}}) \right]_{ir} = \begin{cases} 0 & , \text{ for } \hat{a}_{ir} < 0, \\ \hat{a}_{ir} & , \text{ for } 0 \leq \hat{a}_{ir} \leq 1, \\ 1 & , \text{ for } \hat{a}_{ir} > 1. \end{cases} \quad (31)$$

243 A similar procedure is applied to the updates $\hat{\mathbf{b}}_r$. The updates before projection of its elements for
244 $j' \in \{1, 2, \dots, J\}$ are

$$\hat{b}_{j'r'} = \frac{\sum_{i=1}^I \left[(y_{ij'} - 1 + \hat{q}_{ij'}^{r'}) (\hat{q}_{ij'}^{r'} - \hat{p}_{ij'}^{r'}) \hat{a}_{ir'} \right]}{\sum_{i=1}^I \left[\hat{a}_{ir'} (\hat{q}_{ij'}^{r'} - \hat{p}_{ij'}^{r'}) \right]^2}. \quad (32)$$

245 After executing K updates of all columns of $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$ with PHALS, the elements of $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$
 246 are projected on the set of binary values with $\mathcal{P}_{\mathbb{B}}(\cdot)$ (8).

247 The full algorithm using PHALS to obtain an approximate generalized Boolean factorization is
 248 given in Algorithm 1.

249 3.3. Algorithms for BMF

250 In the specific case of BMF, the most popular Boolean factorization used in practice, the expres-
 251 sions of different quantities of the underlying algorithms can be easily written for any R . In this
 252 subsection, we detail these expressions.

253 By relying on (14), BMF can be written in matrix form as follows:

$$\bar{f}\left(\mathbf{X}^{1:R}\right) = \bigvee_{r=1}^R \mathbf{X}^r = \mathbf{1}_{I \times J} - \boxed{\bullet}_{r=1}^R (\mathbf{1}_{I \times J} - \mathbf{X}^r) \quad (33)$$

254 For given $\mathbf{X}^{1:R}$, the reconstruction error \mathbf{E} can be written as

$$\mathbf{E} = \mathbf{Y} - \mathbf{1}_{I \times J} + \boxed{\bullet}_{r=1}^R (\mathbf{1}_{I \times J} - \mathbf{X}^r) = \mathbf{Y} - \mathbf{1}_{I \times J} + \boxed{\bullet}_{r=1}^R (\mathbf{1}_{I \times J} - \mathbf{a}_r \mathbf{b}_r^\top). \quad (34)$$

255 The $\mathbf{P}_{r'}$ matrices (26) required in GD are

$$\mathbf{P}_{r'} = \boxed{\bullet}_{\substack{s=1 \\ s \neq r'}}^R (\mathbf{1}_{I \times J} - \mathbf{X}^s) = \boxed{\bullet}_{\substack{s=1 \\ s \neq r'}}^R (\mathbf{1}_{I \times J} - \mathbf{a}_s \mathbf{b}_s^\top). \quad (35)$$

256 For PHALS, the quantities that vary depending on the choice of the Boolean function $\bar{f}(\cdot)$ are
 257 $\hat{p}_{ij}^{r'}$ and $\hat{q}_{ij}^{r'}$. For BMF, we have $\hat{p}_{ij}^{r'} = 0$ for all possible tuples (ijr') , while $\hat{q}_{ij}^{r'}$ can be written in
 258 matrix form as $\mathbf{P}_{r'}$ above for $r' \in \{1, \dots, R\}$:

$$\mathbf{Q}_{r'} = \mathbf{P}_{r'} = \boxed{\bullet}_{\substack{s=1 \\ s \neq r'}}^R (\mathbf{1}_{I \times J} - \mathbf{a}_s \mathbf{b}_s^\top). \quad (36)$$

259 4. Numerical experiments

260 In this section, we present the results of numerical experiments concerning the proposed algo-
 261 rithms. We focus first on the BMF setting, that is, when the combining function $f(\cdot)$ is the ‘OR’
 262 function. Under this setting, we compare the performance of the algorithms with 3 other BMF
 263 methods from the literature on simulated noisy binary data. In the first and second simulation

Algorithm 1 Projected hierarchical alternating least squares for general Boolean factorization (PHALS)

Require: \mathbf{Y} , R , K .

- 1: Initialize $\hat{\mathbf{A}}$, $\hat{\mathbf{B}}$ with random i.i.d. uniform elements $\hat{a}_{ir} \sim \mathcal{U}[0, 1]$, $\hat{b}_{ir} \sim \mathcal{U}[0, 1]$, for $i \in \{1, \dots, I\}$, $j \in \{1, \dots, J\}$ and $r \in \{1, \dots, R\}$.
 - 2: **for** $k \in \{1, \dots, K\}$ **do**
 - 3: Update $\hat{\mathbf{A}}$:
 - 4: **for** $r \in \{1, \dots, R\}$ **do**
 - 5: Update each column of $\hat{\mathbf{A}}$:
 - 6: **for** $i \in \{1, \dots, I\}$ **do**
 - 7: Update elements of $\hat{\mathbf{a}}_r$ (30): $\hat{a}_{i'r'} = \frac{\sum_{j=1}^J [(y_{i'j} - 1 + \hat{q}_{i'j}^{r'}) (\hat{q}_{i'j}^{r'} - \hat{p}_{i'j}^{r'}) \hat{b}_{j'r'}]}{\sum_{j=1}^J [\hat{b}_{j'r'} (\hat{q}_{i'j}^{r'} - \hat{p}_{i'j}^{r'})]^2}$
 - 8: **end for**
 - 9: **for** $r' \in \{1, \dots, R\}$ and $r' \neq r$ **do**
 - 10: Update $\hat{p}_{i'j'r'}$ with (28) for $r' \neq r$ and $\hat{q}_{i'j'r'}$ in a similar manner.
 - 11: **end for**
 - 12: **end for**
 - 13: Project onto $[0, 1]$ (31): $\hat{\mathbf{A}} := \mathcal{P}_U(\hat{\mathbf{A}})$
 - 14: Update $\hat{\mathbf{B}}$:
 - 15: **for** $r \in \{1, \dots, R\}$ **do**
 - 16: Update each column of $\hat{\mathbf{B}}$:
 - 17: **for** $i \in \{1, \dots, J\}$ **do**
 - 18: Update elements of $\hat{\mathbf{b}}_r$ (32): $\hat{b}_{j'r} = \frac{\sum_{i=1}^I [(y_{ij'} - 1 + \hat{q}_{ij'}^{r'}) (\hat{q}_{ij'}^{r'} - \hat{p}_{ij'}^{r'}) \hat{a}_{ir'}]}{\sum_{i=1}^I [\hat{a}_{ir'} (\hat{q}_{ij'}^{r'} - \hat{p}_{ij'}^{r'})]^2}$
 - 19: **end for**
 - 20: **for** $r' \in \{1, \dots, R\}$ and $r' \neq r$ **do**
 - 21: Update $\hat{p}_{ij'r'}$ and $\hat{q}_{ij'r'}$.
 - 22: **end for**
 - 23: **end for**
 - 24: Project onto $[0, 1]$ (31): $\hat{\mathbf{B}} := \mathcal{P}_U(\hat{\mathbf{B}})$
 - 25: **end for**
 - 26: Project onto $\{0, 1\}$ (8): $\hat{\mathbf{A}} := \mathcal{P}_{\mathbb{B}}(\hat{\mathbf{A}})$, $\hat{\mathbf{B}} := \mathcal{P}_{\mathbb{B}}(\hat{\mathbf{B}})$
-

264 settings, the data are drawn from random BMF models which are then perturbed by binary flipping
 265 noise. In the first setting, the performance of the algorithms is presented for different number of
 266 columns R of \mathbf{A} and \mathbf{B} , and the probability of binary flipping the data (equivalent to noise intensity)
 267 is kept constant. In the second setting, R is kept constant and results are shown for different values
 268 of the probability of binary flipping. Simulation results will then be presented concerning the sen-
 269 sitivity of the proposed methods to initialization, convergence behavior and time complexity. The
 270 presentation of simulation results on BMF is then followed by its application to real datasets. We
 271 apply PHALS to retrieve the BMF of the following datasets: the congressional voting dataset [34]¹,
 272 the zoo dataset [34]², the New and Old Worlds (NOW) paleontological database [35] and the United
 273 Nations voting dataset [36]³. We end the section by presenting performance results for PHALS and
 274 GD for simulated data generated with $f(\cdot)$ equal to the XOR with two inputs (XOR-2) and to the
 275 3-term majority function (MAJ-3).

276 4.1. Performance for different R

277 In what follows, the performances of the two proposed algorithms PHALS and GD are compared
 278 to 3 state-of-the-art methods for BMF discussed in the introduction: *ASSociation rules* algorithm
 279 (ASSO) [18], the *Formal Concept* (FC) analysis based algorithm [19] and the *Post-NonLinear Penalty*
 280 *Function* (PNL-PF) algorithm [23]. We have also included in the simulations a version of GD that
 281 is initialized with PHALS, we will denote that version of GD as PHALS+GD.

282 In this set of simulations, the data matrix \mathbf{Y} is a 20×20 matrix corresponding to a perturbed
 283 version of a BMF with R components $\mathbf{X} = \bigvee_{r=1}^R \mathbf{a}_r \mathbf{b}_r^\top$. The noise matrix is binary $\mathbf{N} \in \{0, 1\}$
 284 and the perturbation consists in flipping the elements of \mathbf{X} . Therefore, the elements of \mathbf{Y} can be
 285 written using the logical ‘XOR’: $y_{ij} = x_{ij} \oplus n_{ij}$. The elements of \mathbf{N} are drawn iid from a Bernoulli
 286 distribution $n_{ij} \sim \mathcal{B}(p_n)$ where $p_n = \mathbb{P}(n_{ij} = 1)$.

287 For a given data matrix \mathbf{Y} , PHALS and GD are initialized at random $n_{\text{init}} = 3$ times and the
 288 solution achieving the least reconstruction error $\mathcal{F}(\hat{\mathbf{A}}, \hat{\mathbf{B}})$ is kept. For PHALS+GD, GD is initialized
 289 with the best of the 3 initializations of PHALS. PNL-PF is initialized with the result of NMF [22]
 290 applied to the data. The NMF algorithm is initialized randomly as PHALS and GD. ASSO and FC
 291 do not require initializations.

¹<https://archive.ics.uci.edu/ml/datasets/Congressional+Voting+Records>

²<https://archive.ics.uci.edu/ml/datasets/Zoo>

³<https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/LEJUQZ>

292 GD and PNL-PF are executed with $K_{\text{GD/PNL-PF}} = 2000$ iterations for each simulation, while
 293 PHALS is executed $K_{\text{PHALS}} = 1000$ iterations. The step-length of GD is set to a constant $\gamma_k = \gamma =$
 294 0.1 and its hyperparameter λ is increased linearly from 0.01 to 10. The threshold value τ for ASSO
 295 (see [18]) is fixed to 0.9.

296 Under each different simulation setting, $N = 100$ random data matrices \mathbf{Y}^n ($n = 1, \dots, N$)
 297 are generated, each with a random pair $(\mathbf{X}^n, \mathbf{N}^n)$. Each \mathbf{X}^n with a given R is obtained from
 298 randomly generated factor matrices \mathbf{A}^n and \mathbf{B}^n . The elements of \mathbf{A}^n and \mathbf{B}^n are iid samples
 299 from Bernoulli distributions $a_{ir} \sim \mathcal{B}(p_a)$, $b_{ir} \sim \mathcal{B}(p_b)$ where $p_a = p_b = 0.4$. After applying the
 300 algorithms on all \mathbf{Y}^n , their performances in terms of normalized mean square errors (NMSE) of
 301 prediction of \mathbf{X}^n and retrieval of \mathbf{A}^n and \mathbf{B}^n are evaluated. The expressions of these NMSE
 302 are the following: $\text{NMSE}_{\mathbf{X}} = \frac{1}{N I J} \sum_{n=1}^N \left\| \mathbf{X}^n - \bigvee_{r=1}^R \hat{\mathbf{a}}_r^n \left(\hat{\mathbf{b}}_r^n \right)^\top \right\|_F^2$, $\text{NMSE}_{\mathbf{A}} = \frac{1}{N I R} \sum_{n=1}^N \left\| \mathbf{A}^n - \hat{\mathbf{A}}^n \right\|_F^2$,
 303 $\text{NMSE}_{\mathbf{B}} = \frac{1}{N J R} \sum_{n=1}^N \left\| \mathbf{B}^n - \hat{\mathbf{B}}^n \right\|_F^2$. Since the elements of all matrices are binary, these NMSE can
 304 be interpreted as error rates. Note that, due to the permutation ambiguity on the estimation of
 305 the factors, their columns should be permuted to match in the best possible way those of the true
 306 factors before the evaluation of the NMSE.

307 The evolution of the NMSE for the BMF methods is shown in Fig. 1a for $p_n = 0.1$ and $R \in$
 308 $\{2, \dots, 6\}$. As intuitively expected, for most methods, larger values of R lead to larger NMSE both
 309 on prediction of \mathbf{X} and on the retrieved factors. One can also observe that there is no significant
 310 difference in performance between the PHALS, GD and PHALS+GD. PNL-PF has a moderately
 311 inferior performance compared to the proposed methods, while ASSO and FC achieve a significantly
 312 inferior performance in terms of retrieving \mathbf{X} , \mathbf{A} and \mathbf{B} . ASSO and FC do not seem to be adapted
 313 to the approximation setting where noise is present, which has been also observed in previous studies
 314 [23].

315 4.2. Results for different p_n

316 The second simulation setting is very similar to the previously presented one, except that, in
 317 this case, R is kept to a constant value, $R = 3$, and p_n is varied from 0 to 0.3 by increments of 0.02.
 318 Fig. 1b displays the performances of the methods. One can see that, as it is naturally expected, the
 319 performances degrade as p_n increases. As in the previous setting, all the proposed methods seem
 320 to lead to similar performances and they achieve a superior performance compared to the other 3
 321 methods from the literature. Observe also that as p_n gets close to 0.3 all NMSE of the proposed
 322 methods are close to 0.3 and for values smaller than $p_n < 0.3$ the NMSE seem to be smaller than p_n .

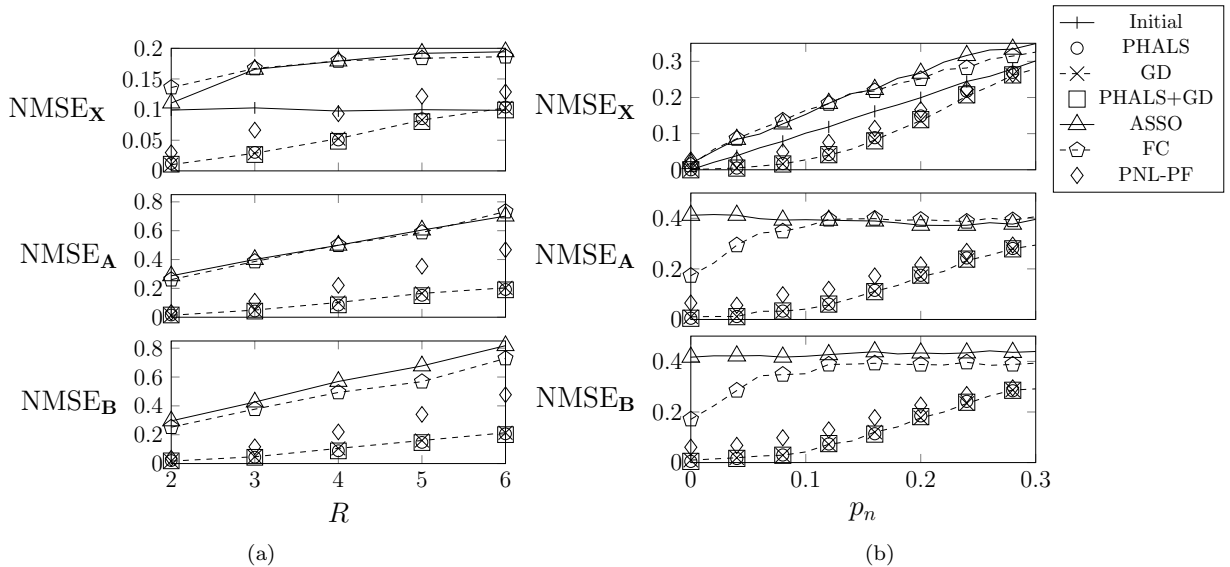


Figure 1: NMSE for the prediction of \mathbf{X} of size 20×20 and retrieval of its BMF factors \mathbf{A} and \mathbf{B} . The factorization algorithms are executed on a noisy version \mathbf{Y} of \mathbf{X} . The noise acts by flipping the elements of \mathbf{Y} with probability p_n . The curves named “Initial” in the top subfigures indicate the NMSE of predicting \mathbf{X} simply using \mathbf{Y} . In (a), the number of columns of matrices \mathbf{A} and \mathbf{B} varies from 2 to 6, while $p_n = 0.1$. In (b), $R = 3$ and p_n is varied from 0 to 0.3 with increments of 0.02.

323 Note, on the top figure, that predicting \mathbf{X} using the noisy data is as efficient as using the proposed
 324 methods for $p_n = 0.3$.

325 Finally, one can see that, when $p_n = 0$, the NMSE on the factors is not zero. This may be due
 326 to convergence of the algorithms to factorizations which are not global minima of (16) or to the non
 327 uniqueness of the approximation of some realizations \mathbf{Y}^n . However, since NMSE are very small for
 328 $p_n = 0$, it seems that the occurrence of such issues is very rare.

329 4.3. Simple example with unique decomposition

330 From the previous simulation results, it seems that ASSO and FC are not adequate in the BMF
 331 approximation setting. Therefore, in what follows, we focus on comparing only PHALS, GD and
 332 PNL-PF.

333 As previously presented, the algorithms may converge to wrong $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$, even in the exact
 334 factorization setting ($p_n = 0$) when the underlying factorization is unique. Due to the non convexity
 335 of the underlying cost functions, not all initializations $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$ lead to the original factors. To try
 336 to assess to which extent the algorithms are prone to this behavior, we have tested them on 3 small

	\mathbf{Y}_1	\mathbf{Y}_4	\mathbf{Y}_5
PHALS	99	100	96
GD	79	100	48
PHALS+GD	99	100	96
PNL-PF	100	66	72

Table 2: Success rate $S_{\%}$ for retrieving the exact BMF for 3 different matrices \mathbf{Y}_1 , \mathbf{Y}_4 and \mathbf{Y}_5 using $n_{\text{init}} = 100$ different random initializations.

337 exact factorization problems with unique factorizations. The 3 considered data matrices are \mathbf{Y}_1 (5),
338 \mathbf{Y}_4 (6) and the following 5×5 matrix from [23]:

$$\mathbf{Y}_5 = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 1 \end{bmatrix}.$$

339 The matrices \mathbf{Y}_1 and \mathbf{Y}_4 have rank 2, while \mathbf{Y}_3 has rank 3. PHALS, GD and PNL-PF are applied
340 to these data with $n_{\text{init}} = 100$ random initializations $\hat{\mathbf{A}}_0^i, \hat{\mathbf{B}}_0^i, i \in \{1, \dots, n_{\text{init}}\}$. The number of
341 iterations for all algorithms is $K = 2000$. Parameters γ_k and λ of GD have been set as in the previous
342 simulations. For each algorithm and each data matrix we have calculated the success rate $S_{\%}$ in
343 percent of retrieving \mathbf{A} and \mathbf{B} from data: $S_{\%} = \left[\text{card} \left(\left\{ i \mid \hat{\mathbf{A}}^i = \mathbf{A} \text{ and } \hat{\mathbf{B}}^i = \mathbf{B} \right\} \right) / n_{\text{init}} \right] \times 100$,
344 where $\text{card}(\cdot)$ denotes the cardinal of a set and $\hat{\mathbf{A}}^i, \hat{\mathbf{B}}^i$ are the output factors of an algorithm when
345 initialized with $\hat{\mathbf{A}}_0^i$ and $\hat{\mathbf{B}}_0^i$. The success rates are displayed in Tab. 2. We observe that the algorithm
346 which seem less prone to converge to spurious factors is PHALS. GD and PNL-PF may converge to
347 spurious factors, but they do not seem to behave equally through the examples. From the results,
348 we can also note that applying GD initialized with the resulting factors from PHALS does not lead
349 to an improvement in the success rate.

350 4.4. Convergence

351 To compare the convergence behavior of PHALS, GD and PNL-PF, we generate 3 random \mathbf{Y}^k
352 in the same manner as in Subsec. 4.1 for R equal to 2, 4 and 6. We then apply the 3 algorithms
353 with $n_{\text{init}} = 100$ and $K = 500$. At each iteration $k \in \{1, \dots, K\}$, we evaluate the overall changes in
354 the factors using the following quantity: $\Delta_k = \frac{\|\hat{\mathbf{A}}_k - \hat{\mathbf{A}}_{k-1}\|_F^2 + \|\hat{\mathbf{B}}_k - \hat{\mathbf{B}}_{k-1}\|_F^2}{\|\hat{\mathbf{A}}_0\|_F^2 + \|\hat{\mathbf{B}}_0\|_F^2}$, where $\hat{\mathbf{A}}_k$ and $\hat{\mathbf{B}}_k$ are the
355 k -th updates of the factors for a given algorithm. The quantity Δ_k is small whenever the factors
356 do not change in consecutive iterations. Therefore, if Δ_k reduces as k increases, the algorithm is
357 converging. Fig. 2 shows some statistics on Δ_k for each algorithm. The statistics displayed are

358 the median, the 5-th and 95-th percentiles of Δ_k evaluated with the $n_{\text{init}} = 100$ available values for
 359 each k . The overall behavior we can observe from this figure is that PHALS is the fastest method
 360 in terms of convergence, while GD is the slowest. One can also note that the 95-th percentile of
 361 Δ_k increases as R increases. Although PHALS is much faster than the other methods to converge,
 362 when R increases, some initializations may lead it to converge slowly or not to converge at all.

363 *Remarks on convergence guarantees:* in their present form, we are not able to give theoretical
 364 guarantees on convergence of the iterates of GD and PHALS.

365 Concerning GD, as presented above, it seems that in practice the algorithm converges, if the
 366 constant step-size $\gamma_k = \gamma$ is chosen sufficiently small. However, theoretical guarantees for convergence
 367 of GD require that a global Lipschitz constant of the gradient of the objective function exists.
 368 Unfortunately, this does not seem to be true for the objective in (16). A possible way to ensure
 369 convergence is to use an adaptive step-length γ_k given by backtracking line-search [37]. With this
 370 modification, since the cost function is analytic and coercive, convergence of GD is guaranteed using
 371 the results from [38].

372 PHALS is a block coordinate descent algorithm. For this class of algorithms, convergence of the
 373 iterates can be ensured, for example, using the results of [39]. To use the results of [39], the objective
 374 function should be separately strongly convex in each block of variables $\mathbf{a}_1, \dots, \mathbf{a}_R, \mathbf{b}_1, \dots, \mathbf{b}_R$.
 375 Unfortunately, this cannot be guaranteed, and, in practice, one can see that for difficult factorization
 376 cases (large R), some initializations may lead to non-converging iterates. One possibility to solve
 377 this issue is to add proximal terms to the objective function at each update. For the updates of \mathbf{a}_r ,
 378 one should add $\rho \|\mathbf{a}_r - \hat{\mathbf{a}}_r\|_2^2$, where $\rho > 0$ is a pre-defined constant and $\hat{\mathbf{a}}_r$ is the most recent update
 379 of \mathbf{a}_r . Similarly, for the updates \mathbf{b}_r , the term $\rho \|\mathbf{b}_r - \hat{\mathbf{b}}_r\|_2^2$ should be added. With modified updates
 380 considering this additional terms, it may be possible to use the results of [39] to ensure convergence
 381 of the iterates.

382 4.5. Time complexity

383 Using a similar simulation setting from the previous subsection, we have also measured the
 384 execution time for the 3 algorithms to finish 2000 iterations. The statistics on execution time⁴ for
 385 100 runs of the algorithms and for $R \in \{2, \dots, 6\}$ are shown in Tab. 3. We observe that PNL-PF
 386 takes much less time than the other methods. GD is from 10 to 40 times slower than PNL-PF and we

⁴These simulations are realized in *Scilab* version 6.1.0 with a processor Intel®Core™ i7-7820HQ, 2.90GHz and with 16GB of RAM.

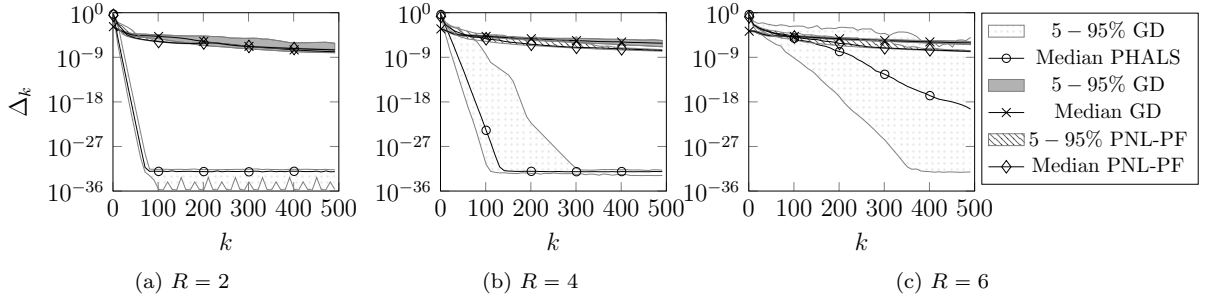


Figure 2: Statistics about the overall changes Δ_k (4.4) in the updates $\hat{\mathbf{A}}_k$ and $\hat{\mathbf{B}}_k$ for GD, PHALS and PNL-PF for the approximate factorization of 3 noisy binary matrices of size 20×20 . Each matrix is generated by randomly drawing a binary BMF model with a given R , then applying binary flipping noise. The subfigures are generated with different values of R , they are indicated in the subcaptions. The statistics are evaluated for $n_{\text{init}} = 100$ different random initializations. The extremes of the bands (5% – 95%) around the median are the 5-th and 95-th percentiles of Δ_k .

		$R = 2$	3	4	5	6
PHALS	5-th percentile	0.844	1.547	2.188	3.0313	4.281
	Median	0.703	1.297	2.000	2.906	3.938
	95-th percentile	0.610	1.219	1.813	2.781	3.750
GD	5-th percentile	1.734	2.734	3.438	24.438	32.688
	Median	1.547	2.359	3.281	23.422	30.500
	95-th percentile	1.406	2.219	3.141	22.313	29.125
PNL-PF	5-th percentile	0.203	0.250	0.234	1.000	0.984
	Median	0.156	0.172	0.172	0.672	0.719
	95-th percentile	0.094	0.125	0.125	0.422	0.469

Table 3: Statistics on total execution times in seconds for approximate R -component BMF of 20×20 binary matrices. The statistics are evaluated for 100 runs of the 3 algorithms with $K = 2000$ iterations in each run.

387 can clearly see a large relative increase in execution for GD when passing from $R = 4$ to $R = 5$. Such
388 an increase can also be observed in a lesser extent in PNL-PF, while in PHALS the relative increase
389 is smaller than a factor of 2. Since real datasets may be of sizes much larger than 20×20 , it is clear
390 from these simulations that GD cannot be reasonably used in practical data analysis problems with
391 the implementation used in this work.

392 4.6. Discussion on the results

393 In terms of approximation performance PHALS leads to better results than PNL-PF at the
394 expense of a longer execution time per iteration and of a risk of producing slowly or non converging

395 iterations for large values of R (in our simulations mainly for $R \geq 6$). Note however that for small
396 R , Δ_k for PHALS decreases much faster than for PNL-PF, thus if a threshold on Δ_k is used as
397 convergence criterion to stop the algorithm, the longer execution times of PHALS iterations are
398 compensated by its much faster convergence.

399 4.7. Real datasets

400 In what follows, we obtain the approximate BMF of different real binary datasets. From the
401 previous results on factorization performance obtained through simulations, factorization results
402 are expected to be mostly similar for PHALS, GD and PNL-PF. Therefore, we have applied only
403 PHALS to factorize the real datasets. For each of the datasets, $n_{\text{init}} = 10$ random initializations are
404 used and the best solution in terms of data reconstruction error is selected. The maximum allowed
405 number of iterations is set to $K_{\text{PHALS}} = 2000$ and a convergence criterion based on Δ_k is used as
406 an additional stopping criterion.

407 *US Congressional voting dataset.* We first apply PHALS to a dataset containing 16 key votes of the
408 United States congress for the year 1984 [34]⁵. The votes of 435 representatives are coded by binary
409 values: ‘1’ for a vote in favor of the proposed bill and ‘0’ for a vote against it. Missing votes in a
410 given bill have been replaced by the corresponding majority vote. This allows to fully encode the
411 dataset into a binary matrix of size 435×16 . The dataset also contains information on the party
412 of each representative (democrat or republican). Since there are 2 parties, PHALS is applied to the
413 dataset with $R = 2$. The dataset plot and an illustration of the results are given in Fig. 3 (a-d).
414 One can clearly observe that the patterns related to each component have almost disjoint support,
415 indicating an opposing voting pattern for each component. The error rate on the reconstructed
416 data using the retrieved BMF model is of 20%. By comparing the grouping of the representatives
417 encoded by matrix $\hat{\mathbf{A}}$ with the information on the parties of each candidate, we found that PHALS
418 can predict the party of the representative with an accuracy of 77%.

419 *Zoo dataset.* We also applied PHALS to a dataset containing the information on the presence or
420 absence of a given feature, for example *hair*, *feathers*, *milk*, for different animals. The dataset [34]⁶
421 contains 15 binary features and an integer feature with the number of legs. These features are given
422 for 101 animals. The animals in the dataset are categorized in 7 classes: mammals, birds, reptiles,

⁵<https://archive.ics.uci.edu/ml/datasets/Congressional+Voting+Records>

⁶<https://archive.ics.uci.edu/ml/datasets/Zoo>

423 fishes, amphibians, insects and a class containing many different invertebrate animals (*e.g.* *crab*,
424 *worm*, *octopus*). We have encoded the integer variable corresponding to the number of legs using
425 one-hot encoding. We apply PHALS to the resulting 101×19 binary matrix to try to group the
426 animals by looking at the patterns given by the columns of $\hat{\mathbf{A}}$. The algorithm is applied with R in
427 the range 2 to 7 and the data reconstruction error for these values of R are respectively 0.170, 0.116,
428 0.0928, 0.0771, 0.0693 and 0.0620. One can clearly see that beyond $R = 3$ the improvement on data
429 reconstruction obtained by increasing R is mild. This result seems to be similar to what has been
430 presented in [23] for the analysis of the same dataset with PNL-PF.

431 The dataset and an illustration of the results for $R = 3$ are given in Fig. 3 (f-i). To simplify
432 the interpretation of the results, the rows of the matrices, which correspond to different animals,
433 have been reordered to correspond to continuous blocks of animals of the same category. Reordering
434 has been carried out using the same order of the classes mentioned above, thus the first block of
435 animals correspond to mammals, the second to birds, *etc.* One can observe that component 1 clearly
436 corresponds to a continuous block of animals, in these case mammals. The second component mostly
437 group together birds with two exceptions, *fruitbat* and *vampire*, which are also present in the group
438 of mammals. The third group contains mostly fishes, but also some mammals (*e.g.* *dolphin*) and
439 birds (*e.g.* *penguin*). Many insects and animals from the last category of invertebrate animals are
440 not contained in any components. As the number of components is increased to $R = 7$, it has been
441 observed that the retrieved BMF is not able to clearly separate the 7 underlying categories.

442 *Paleontological dataset.* Following closely [15, 40], we analyze data containing information on the
443 localization of fossil mammals [35]. The objective is to apply PHALS to factorize a binary data
444 matrix where the rows correspond to different genera of fossil mammals and the columns correspond
445 to the different localities where they have been found. The data obtained from [35] is preprocessed
446 in a similar manner as in [15, 40]. Fossils of small mammals are excluded from the dataset and
447 only those retrieved in Europe are considered. We also removed genera which are too infrequent
448 (less than 10 occurrences) and localities where only 1 genera has been found. As a result of this
449 preprocessing, a binary matrix of size 254 (genera) \times 1375 (localities) is obtained, where a ‘1’ stands
450 for the occurrence of at least one fossil of a given genus in a given locality and a ‘0’ for its absence.
451 We have applied PHALS to this dataset to see if the resulting BMF allows to find communities of
452 mammals that appear in similar localities. The algorithm has been applied with R in the range
453 2 – 7. The algorithm seems to suffer from convergence issues for $R > 4$, generating factor matrices
454 with spurious empty columns. To validate the results for $R \leq 4$, we have followed [40] and plotted

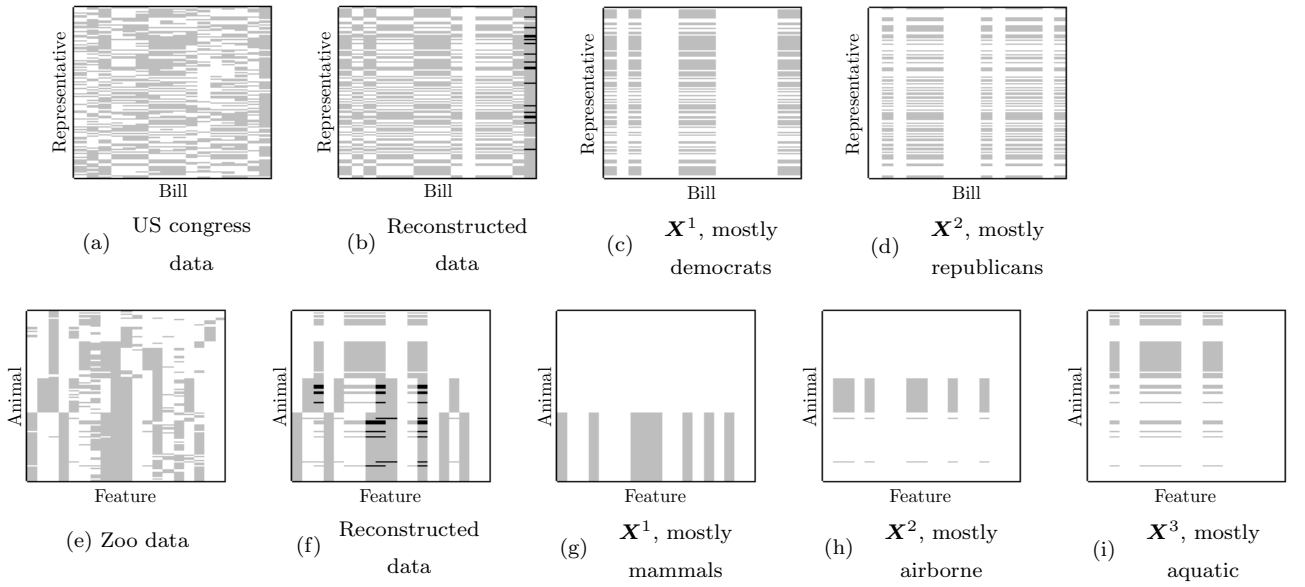


Figure 3: Real datasets and the results obtained with PHALS. In (a) and (e) the US congress voting dataset and the Zoo dataset are displayed. The gray color corresponds to a ‘1’ in the underlying dataset matrix, while white color corresponds to ‘0’. In (b) and (f), the reconstructed data using the BMF model are displayed. The black color indicates intersections between BMF components’ supports. The number of BMF components are $R = 2$ and $R = 3$ respectively. The rank-1 components \mathbf{X}^r retrieved with PHALS are displayed on the right of the reconstructed data in (c), (d), (g), (h) and (i).




Figure 4: Minimum age in millions of years [myr] related to the genera of the different groups obtained by applying PHALS with $R = 4$ to the paleontology dataset [35].

455 the values of a variable related to the minimum age in millions of years of the localities where the
 456 different genera have been found. We have observed that as R increases PHALS finds groups of
 457 fossils with increasing minimum age. The ages of the genera in the different groups for $R = 4$ are
 458 displayed in Fig. 4. The genera in each of the components of this figure are the following:

- 459 • Component 1: *Amphiperatherium*, *Amphitragulus*, *Andegameryx*, *Brachyodus*, *Cainotherium*,
 460 *Cynelos*, *Diaceratherium*, *Palaeogale*, *Protaceratherium*.
- 461 • Component 2: *Amphicyon*, *Anchitherium*, *Anisodon*, *Aureliachoerus*, *Brachypotherium*, *Buno-*
 462 *listriodon*, *Dicrocerus*, *Dorcatherium*, *Gomphotherium*, *Hemicyon*, *Hyotherium*, *Lagomeryx*,
 463 *Lartetotherium*, *Listriodon*, *Martes*, *Micromeryx*, *Palaeomeryx*, *Plesiaceratherium*, *Procervu-*
 464 *lus*, *Prodeinotherium*, *Prosantorhinus*, *Pseudaelurus*, *Styriofelis*, *Taucanamo*.
- 465 • Component 3: *Adcrocuta*, *Choerolophodon*, *Cremohipparion*, *Deinotherium*, *Dihoplus*, *Gazella*,
 466 *Helladotherium*, *Hipparion*, *Hippopotamodon*, *Hippotherium*, *Hyaenictitherium*, *Miotragocerus*,
 467 *Palaeotragus*, *Pliodiceros*, *Tragoportax*.
- 468 • Component 4: *Bison*, *Canis*, *Cervus*, *Equus*, *Lynx*, *Mammuthus*, *Panthera*, *Stephanorhinus*,
 469 *Sus*, *Ursus*, *Vulpes*.

470 By inspecting the median minimum ages of these groups, it seems that PHALS is able to retrieve
 471 animal communities that have lived in different ages.

472 *UN voting dataset*. As a last example of application, we analyze the grouping of countries produced
 473 by PHALS when used to factorize a binary matrix generated from the UN voting dataset [36]⁷. We

⁷<https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/LEJUQZ>

474 follow a similar setting as considered in [15] and we encode in a binary matrix the votes during the
475 cold-war period (1946 – 1990) of different countries for different UN roll-calls. We have removed
476 from the dataset all roll-calls whose number of unknown votes is larger than 98 (half of the listed
477 countries) and also all the roll-calls with unanimous results. The unknown votes in the remaining
478 roll-calls have been replaced by the majority vote. For this dataset, encoding with ‘1’ votes in favor
479 of a roll-call leads to a very dense data matrix, whose BMF is difficult to retrieve and interpret.
480 Therefore, to have a more sparse data matrix, we have encoded with ‘1’ votes against a roll-call and
481 with ‘0’, votes in favor of it. PHALS has been applied to this dataset with R in the range 2 – 7. The
482 data reconstruction error is respectively 0.0312, 0.0241, 0.0215, 0.0194, 0.0173 and 0.0158. Although
483 a large part of the approximation improvement is observed when increasing R from 2 to 3, when we
484 analyze the groups of countries produced for each R , interesting results seem to appear up to $R = 6$.

485 When $R = 2$, we can find a component containing the following countries: Australia, Belgium,
486 Canada, Denmark, France, West Germany, Iceland, Israel, Italy, Japan, Luxembourg, Netherlands,
487 New Zealand, Norway, Portugal, UK, US. The second component contains 173 countries from dif-
488 ferent continents. If R is increased to 3, the first 2 components are similar to those obtained with
489 $R = 2$ and a third component groups countries from the socialist block: Belarus, Bulgaria, Cuba,
490 Czechoslovakia, East Germany, Hungary, Mongolia, Poland, Russia, Ukraine. When increasing R
491 to 4, similar results are obtained and a component with US and Israel appears. While for $R = 5$,
492 2 components with countries from the occidental block of countries are produced. For $R = 6$, the
493 component containing a large number of countries seems to contain much less countries than for
494 smaller R and a sixth component containing 36 countries from different continents appear. This
495 last component gathers countries from the previously obtained component with a large number of
496 countries but also countries from the occidental block (*e.g.* France) and from the socialist block (*e.g.*
497 Cuba). Finally, for $R = 7$, the algorithm start finding components containing single countries (*e.g.*
498 a component with only US).

499 4.8. Results for XOR – 2 and MAJ – 3

500 The last experimental results concern the application of the GD and PHALS to a factorization
501 setting different from BMF. We consider two other Boolean combining functions, the logical ‘XOR’
502 with two inputs $x_1 \oplus x_2$ (XOR – 2) and 3-term majority $\mathbb{1}_{(\sum_i x_i) \geq 2}(x_1, x_2, x_3)$ (MAJ – 3). Since the
503 uniqueness properties of these factorizations are still very little understood, we only focus on testing
504 the methods for data denoising. We consider a simulation setting similar to the one presented in
505 Subsection 4.2, the main differences are that the underlying (clean) data are generated with the




Figure 5: NMSE for the prediction of \mathbf{X} of size 20×20 for generalized Boolean factorizations using the XOR – 2 and MAJ – 3 component combining functions. The decomposition algorithms are executed on a noisy version \mathbf{Y} of \mathbf{X} . The probability p_n of the binary noise that flips the elements of \mathbf{X} is varied from 0 to 0.3 with increments of 0.02.

506 XOR and MAJ-3 functions, the maximum allowed number of iterations of the algorithms is set to
 507 $n_{\text{init}} = 5000$ and that an additional stopping criterion based Δ_k is used for ending the iterations.
 508 The results for the NMSE of data reconstruction are displayed in Fig. 5. One can observe that
 509 the algorithms denoise the data, since their NMSE is smaller than the NMSE for the noisy data
 510 (curve named *Initial* in the plot). In both cases, GD seems to achieve a slightly superior denoising
 511 performance than PHALS. It is also possible to observe that the factorizations do not seem to have
 512 the same robustness behavior against noise. The denoising performance for F2MF (XOR-2) for small
 513 noise intensity seems much superior than for the MAJ-3 factorization. This is intuitively expected,
 514 since the MAJ-3 factorization requires the estimation of more parameters for the same amount of
 515 data. For large noise intensities ($p_n \approx 0.3$), the opposite behavior is observed, with the MAJ-3
 516 factorization leading to a superior denoising performance.

517 5. Conclusions and further work

518 In this paper, we have introduced a generalized framework for the Boolean factorization of binary
 519 matrices, where the “sum” between the rank-1 binary terms can be an arbitrary Boolean function.
 520 We proposed two iterative algorithms for achieving this factorization, based on gradient descent
 521 (GD) and on projected hierarchical alternating least squares (PHALS) approaches, respectively.
 522 Implementation details for the algorithms have been presented for BMF and compared through
 523 numerical experiments with state-of-the art algorithms from the literature.

524 From the results of the numerical experiments, it seems that the best performing algorithm is

525 PHALS, both in terms of performance of retrieving the factorization and of overall computation
526 time. Although GD gives good results in terms of approximate factorization performance, its high
527 complexity impedes its practical use on large datasets.

528 We have also tested PHALS to retrieve the BMF of real datasets. The components obtained
529 seem to agree with those obtained in other works of the literature and with intuition on what would
530 be possible groupings of the data. In this paper, we have not focused on the choice of the number
531 of components R and in the presentation of the results for the real datasets, we have chosen a value
532 of R that seemed to give stable results with components agreeing with intuition on the dataset. In
533 practice, if no intuition on the expected components is available, a quantitative criterion for choosing
534 R may be used. Such criteria will be studied and tested in future work.

535 At the end of the experimental section, we have also presented results of applying PHALS and
536 GD in a more general factorization setting where XOR – 2 and MAJ – 3 component combining
537 functions are considered instead of the logical OR of BMF. Such matrices factorizations are not
538 identifiable in general and thus may not be useful in data analysis. In future work, we would like to
539 extend our general approach to the higher-order tensor setting and to verify whether the extended
540 models are identifiable.

541 **References**

- 542 [1] T. Li, A general model for clustering binary data, in: Proceedings of the eleventh ACM SIGKDD
543 International Conference on Knowledge Discovery in Data Mining, 2005, pp. 188–197.
- 544 [2] L. Kozma, A. Ilin, T. Raiko, Binary principal component analysis in the Netflix collaborative
545 filtering task, in: 2009 IEEE International Workshop on Machine Learning for Signal Processing,
546 IEEE, 2009, pp. 1–6.
- 547 [3] E. Nenova, D. I. Ignatov, A. V. Konstantinov, An FCA-based boolean matrix factorisation for
548 collaborative filtering, in: FCAIR 2012 Formal Concept Analysis Meets Information Retrieval
549 Workshop co-located with the 35th European Conference on Information Retrieval (ECIR 2013)
550 March 24, 2013, Moscow, Russia, 2013, p. 57.
- 551 [4] M. Diop, S. Miron, A. Larue, D. Brie, Binary matrix factorization applied to Netflix dataset
552 analysis, IFAC-PapersOnLine 52 (24) (2019) 13–17.
- 553 [5] E. Meeds, Z. Ghahramani, R. M. Neal, S. T. Roweis, Modeling dyadic data with binary latent
554 factors, Advances in neural information processing systems 19 (2007) 977.

- 555 [6] Z.-Y. Zhang, T. Li, C. Ding, X.-W. Ren, X.-S. Zhang, Binary matrix factorization for analyzing
556 gene expression data, *Data Mining and Knowledge Discovery* 20 (1) (2010) 28–52.
- 557 [7] S. Tu, L. Xu, R. Chen, A binary matrix factorization algorithm for protein complex predic-
558 tion, in: 2010 IEEE International Conference on Bioinformatics and Biomedicine Workshops
559 (BIBMW), IEEE, 2010, pp. 113–118.
- 560 [8] H. Lu, J. Vaidya, V. Atluri, Optimal Boolean matrix decomposition: Application to role engi-
561 neering, in: 2008 IEEE 24th International Conference on Data Engineering, IEEE, 2008, pp.
562 297–306.
- 563 [9] S. Talwar, M. Viberg, A. Paulraj, Blind separation of synchronous co-channel digital signals
564 using an antenna array. I. algorithms, *IEEE Transactions on Signal Processing* 44 (5) (1996)
565 1184–1197.
- 566 [10] A.-J. Van der Veen, Analytical method for blind binary signal separation, *IEEE Transactions*
567 *on Signal Processing* 45 (4) (1997) 1078–1082.
- 568 [11] A. I. Schein, L. K. Saul, L. H. Ungar, A generalized linear model for principal component
569 analysis of binary data, in: *International Workshop on Artificial Intelligence and Statistics*,
570 PMLR, 2003, pp. 240–247.
- 571 [12] J. De Leeuw, Principal component analysis of binary data by iterated singular value decompo-
572 sition, *Computational statistics & data analysis* 50 (1) (2006) 21–39.
- 573 [13] S. Lee, J. Z. Huang, J. Hu, Sparse logistic principal components analysis for binary data, *The*
574 *annals of applied statistics* 4 (3) (2010) 1579.
- 575 [14] Z. Kang, C. J. Spanos, Sequential logistic principal component analysis (SLPCA): Dimensional
576 reduction in streaming multivariate binary-state system, in: 2014 13th International Conference
577 on Machine Learning and Applications, IEEE, 2014, pp. 171–177.
- 578 [15] A. Lumbreras, L. Filstroff, C. Févotte, Bayesian mean-parameterized nonnegative binary matrix
579 factorization, *Data Mining and Knowledge Discovery* 34 (6) (2020) 1898–1935.
- 580 [16] Z. Zhang, T. Li, C. Ding, X. Zhang, Binary matrix factorization with applications, in: *Seventh*
581 *IEEE International Conference on Data Mining (ICDM 2007)*, IEEE, 2007, pp. 391–400.

- 582 [17] N. Gillis, S. A. Vavasis, On the complexity of robust PCA and ℓ_1 -norm low-rank matrix ap-
583 proximation, *Mathematics of Operations Research* 43 (4) (2018) 1072–1084.
- 584 [18] P. Miettinen, T. Mielikäinen, A. Gionis, G. Das, H. Mannila, The discrete basis problem, *IEEE*
585 *transactions on knowledge and data engineering* 20 (10) (2008) 1348–1362.
- 586 [19] R. Belohlavek, V. Vychodil, Discovery of optimal factors in binary data via a novel method of
587 matrix decomposition, *Journal of Computer and System Sciences* 76 (1) (2010) 3–20.
- 588 [20] M. Trnecka, R. Vyjidacek, Revisiting the Grecon algorithm for Boolean matrix factorization,
589 *Knowledge-Based Systems* (2022) 108895.
- 590 [21] T. Makhalova, M. Trnecka, From-below Boolean matrix factorization algorithm based on MDL,
591 *Advances in Data Analysis and Classification* 15 (1) (2021) 37–56.
- 592 [22] D. D. Lee, H. S. Seung, Learning the parts of objects by non-negative matrix factorization,
593 *Nature* 401 (6755) (1999) 788–791.
- 594 [23] S. Miron, M. Diop, A. Larue, E. Robin, D. Brie, Boolean decomposition of binary matrices
595 using a post-nonlinear mixture approach, *Signal Processing* 178 (2021) 107809.
- 596 [24] D. DeSantis, E. Skau, D. P. Truong, B. Alexandrov, Factorization of binary matrices: Rank
597 relations, uniqueness and model selection of Boolean decomposition, *ACM Transactions on*
598 *Knowledge Discovery from Data (TKDD)* (2020).
- 599 [25] H. Nguyen, R. Zheng, Binary independent component analysis with or mixtures, *IEEE Trans-*
600 *actions on Signal Processing* 59 (7) (2011) 3168–3181.
- 601 [26] S. Ravanbakhsh, B. Póczos, R. Greiner, Boolean matrix factorization and noisy completion via
602 message passing, in: *International Conference on Machine Learning*, PMLR, 2016, pp. 945–954.
- 603 [27] A. Cichocki, A.-H. Phan, Fast local algorithms for large scale nonnegative matrix and tensor fac-
604 torizations, *IEICE transactions on fundamentals of electronics, communications and computer*
605 *sciences* 92 (3) (2009) 708–721.
- 606 [28] D. DeSantis, E. Skau, B. Alexandrov, Factorizations of binary matrices–rank relations and the
607 uniqueness of Boolean decompositions, *arXiv preprint arXiv:2012.10496* (2020).
- 608 [29] J. E. Cohen, U. G. Rothblum, Nonnegative ranks, decompositions, and factorizations of non-
609 negative matrices, *Linear Algebra and its Applications* 190 (1993) 149–168.

- 610 [30] K. H. Kim, Boolean matrix theory and applications, Vol. 70, Dekker, 1982.
- 611 [31] V. L. Watts, Boolean rank of Kronecker products, *Linear Algebra and its Applications* 336 (1-3)
612 (2001) 261–264.
- 613 [32] T. Watson, Nonnegative rank vs. binary rank, arXiv preprint arXiv:1603.07779 (2016).
- 614 [33] Y. Crama, P. L. Hammer, Boolean functions: Theory, algorithms, and applications, Cambridge
615 University Press, 2011.
- 616 [34] D. Dua, C. Graff, UCI machine learning repository (2017).
617 URL <http://archive.ics.uci.edu/ml>
- 618 [35] The now community. new and old worlds database of fossil mammals (now)., [https://](https://nowdatabase.org/now/database/)
619 nowdatabase.org/now/database/, accessed: 2022-06-15.
- 620 [36] E. Voeten, Data and analyses of voting in the un general assembly, in: B. Reinalda (Ed.), *Data*
621 *and Analyses of Voting in the UN General Assembly*, Routledge Londres, 2013.
- 622 [37] J. Nocedal, S. Wright, *Numerical optimization*, Springer Science & Business Media, 2006.
- 623 [38] P.-A. Absil, R. Mahony, B. Andrews, Convergence of the iterates of descent methods for analytic
624 cost functions, *SIAM Journal on Optimization* 16 (2) (2005) 531–547.
- 625 [39] Y. Xu, W. Yin, A block coordinate descent method for regularized multiconvex optimization
626 with applications to nonnegative tensor factorization and completion, *SIAM Journal on imaging*
627 *sciences* 6 (3) (2013) 1758–1789.
- 628 [40] E. Bingham, A. Kabán, M. Fortelius, The aspect Bernoulli model: multiple causes of presences
629 and absences, *Pattern Analysis and Applications* 12 (1) (2009) 55–78.